Privacy-Preserving Hyperparameter Tuning for Federated Learning

Apostolos Pyrgelis

Senior Researcher @ RISE Research Institutes of Sweden

AI & Security Webinar - 30/4/2025

(Cross-Silo) Federated Learning

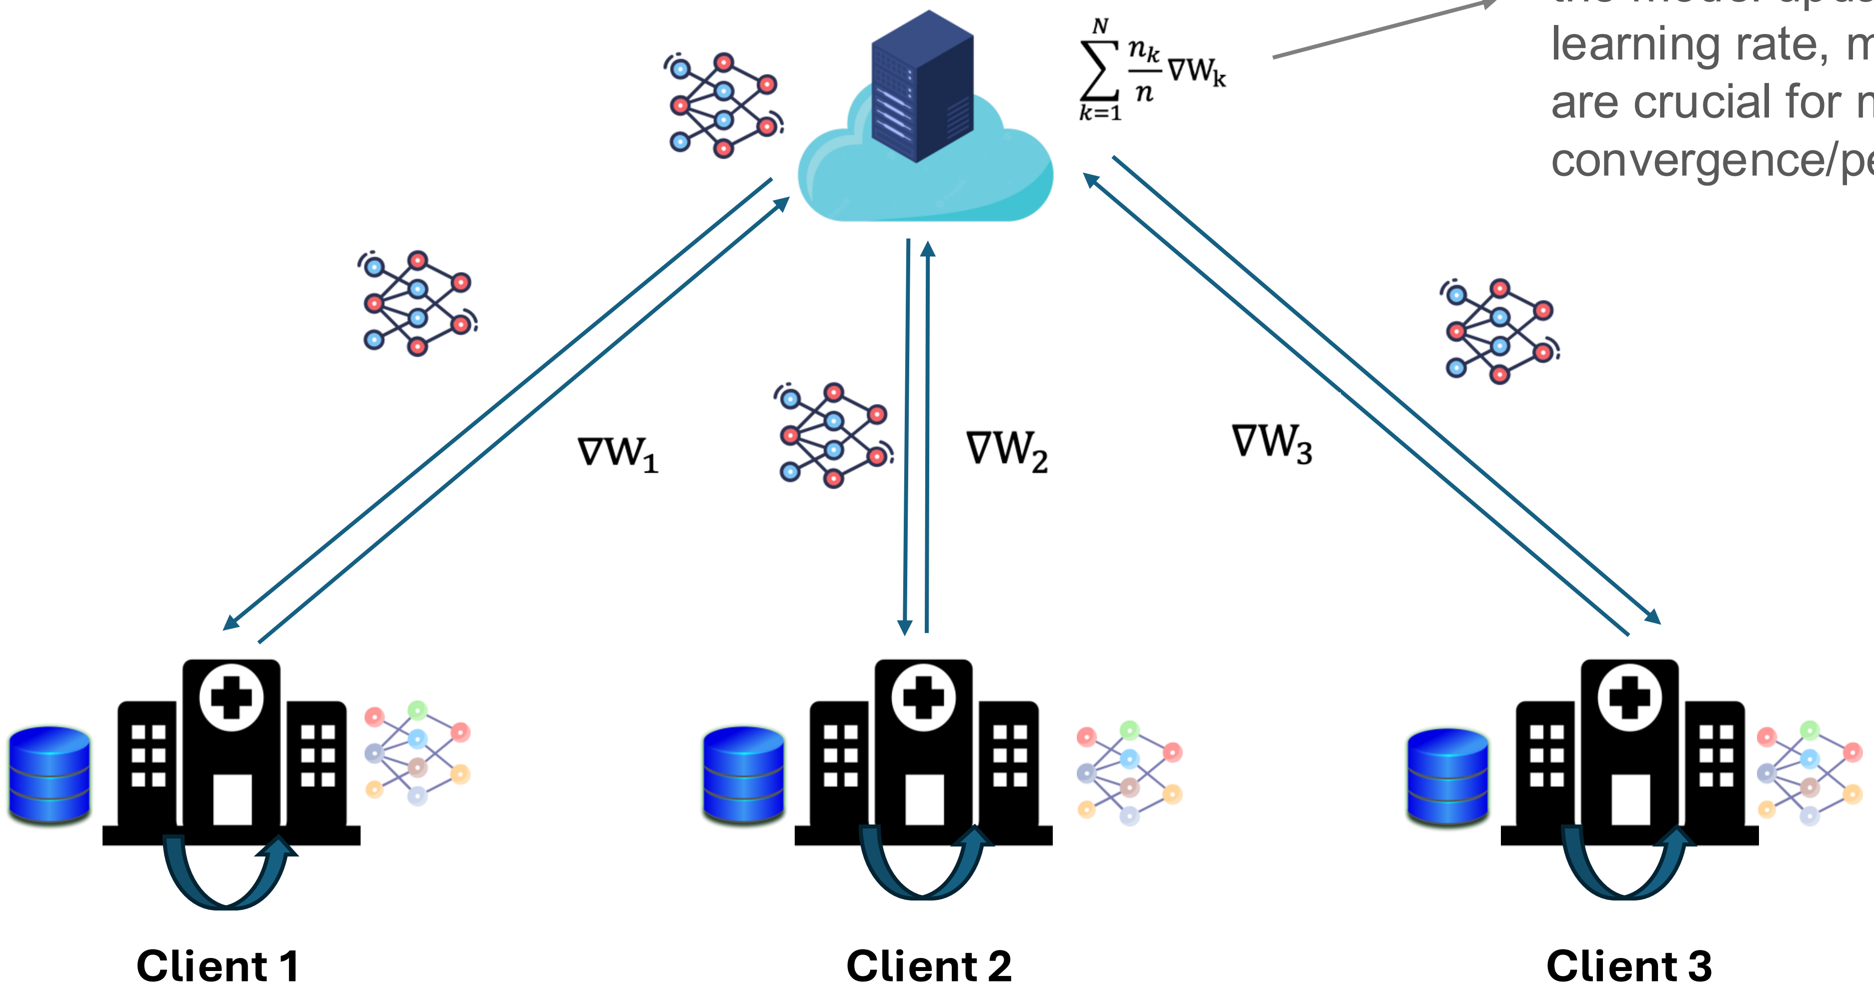

Global hyperparameters for the model update, e.g., learning rate, momentum, are crucial for model convergence/performance

Privacy in Federated Learning (?)

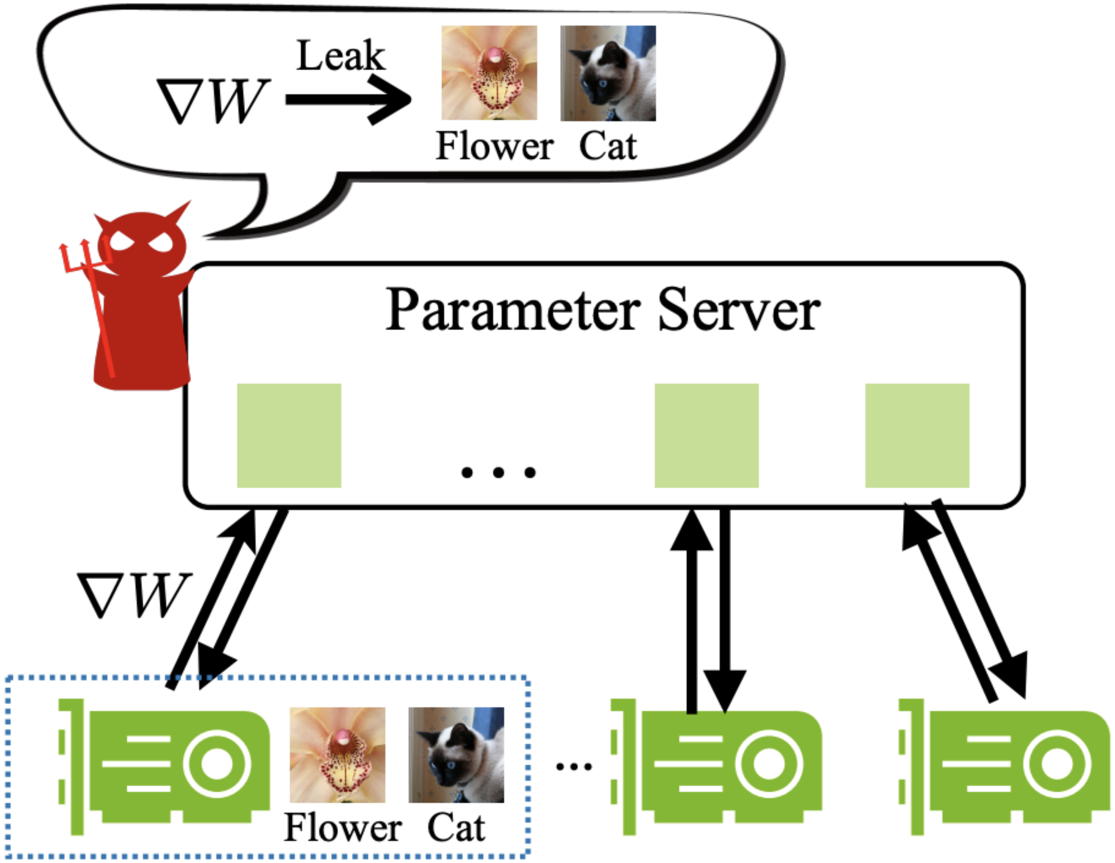

Figure from [1]

[1] Zhu, Ligeng, et al., "Deep leakage from gradients." in *NeurIPS*, 2019.

Privacy-Preserving Federated Learning (PPFL)

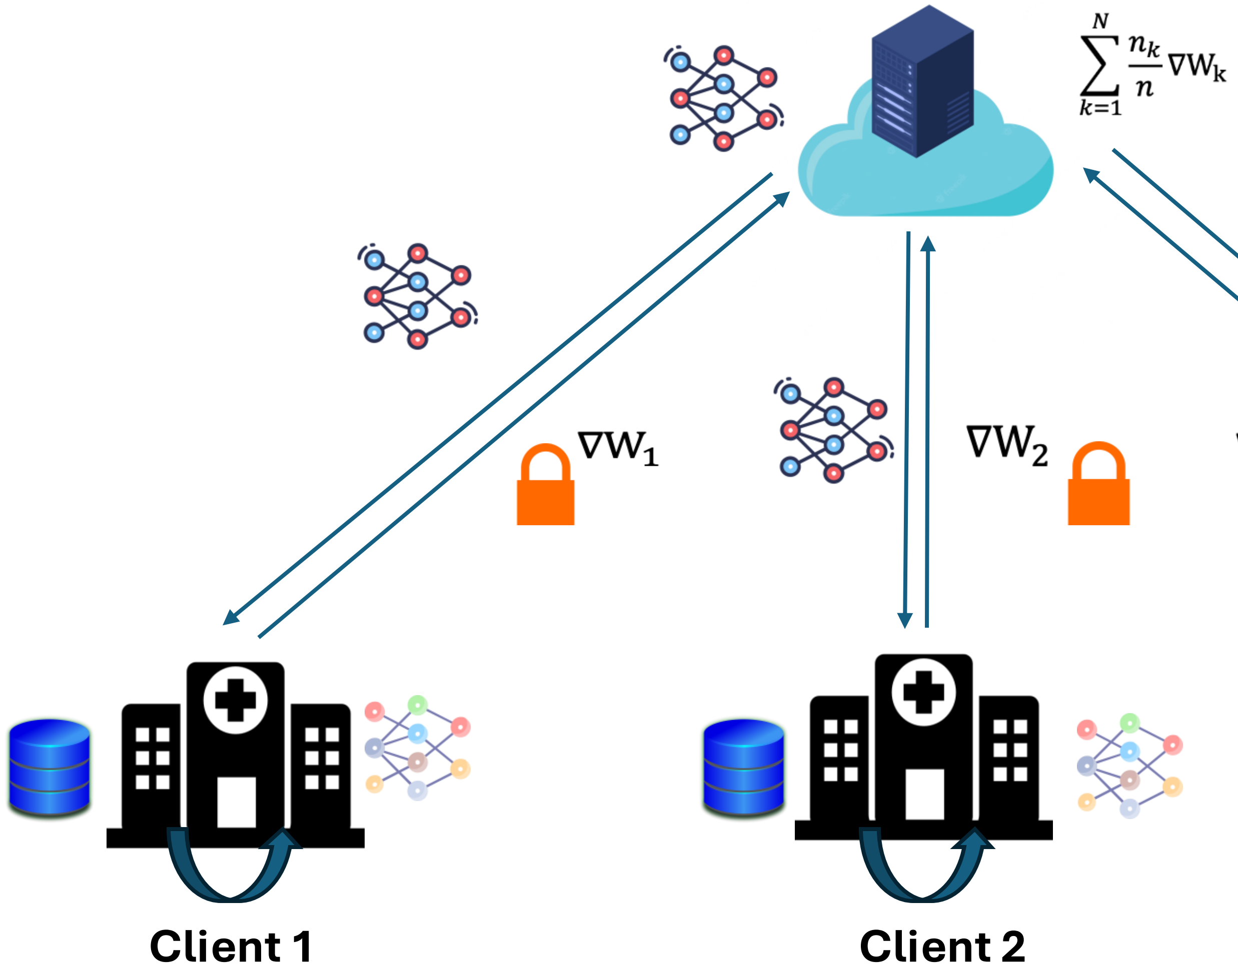

Privacy-enhancing Technologies (PETS):

- Secure Multiparty Computation (MPC)[1]
- Homomorphic Encryption (HE)[2]
- Differential Privacy (DP)[3]

[1] Wagh et al., Falcon: Honest-Majority Maliciously Secure Framework for Private Deep Learning, in *PETS*, 2021.
[2] Sav et al., POSEIDON: Privacy-Preserving Federated Neural Network Learning, in *NDSS*, 2021.
[3] Abadi et al., Deep Learning with Differential Privacy, in *ACM CCS*, 2016.

Hyperparameter Tuning in (PP)FL

- Traditional HP tuning methods, e.g., grid/random search, are impractical for FL settings
- PETS for PPFL raise additional challenges:
  - Computation/communication overhead (HE/MPC)
  - Spending privacy budget on HP tuning (DP)
- Method requirements:
  - Efficiency (single-shot, before the FL training starts)
  - Accuracy (retain the performance of the trained model)
  - Privacy (HP tuning should not result to new forms of leakage)[1]

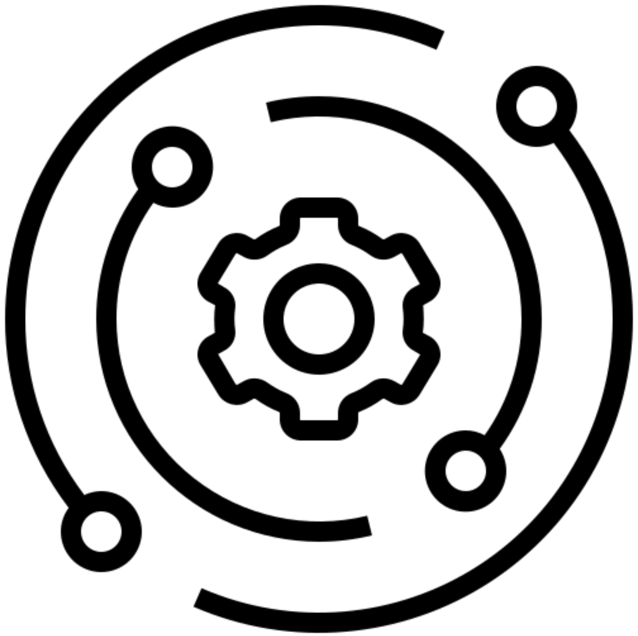

[1] N. Papernot and T. Steinke, "Hyperparameter tuning with Renyi differential privacy," in *ICLR*, 2022.

Overview of our Approach

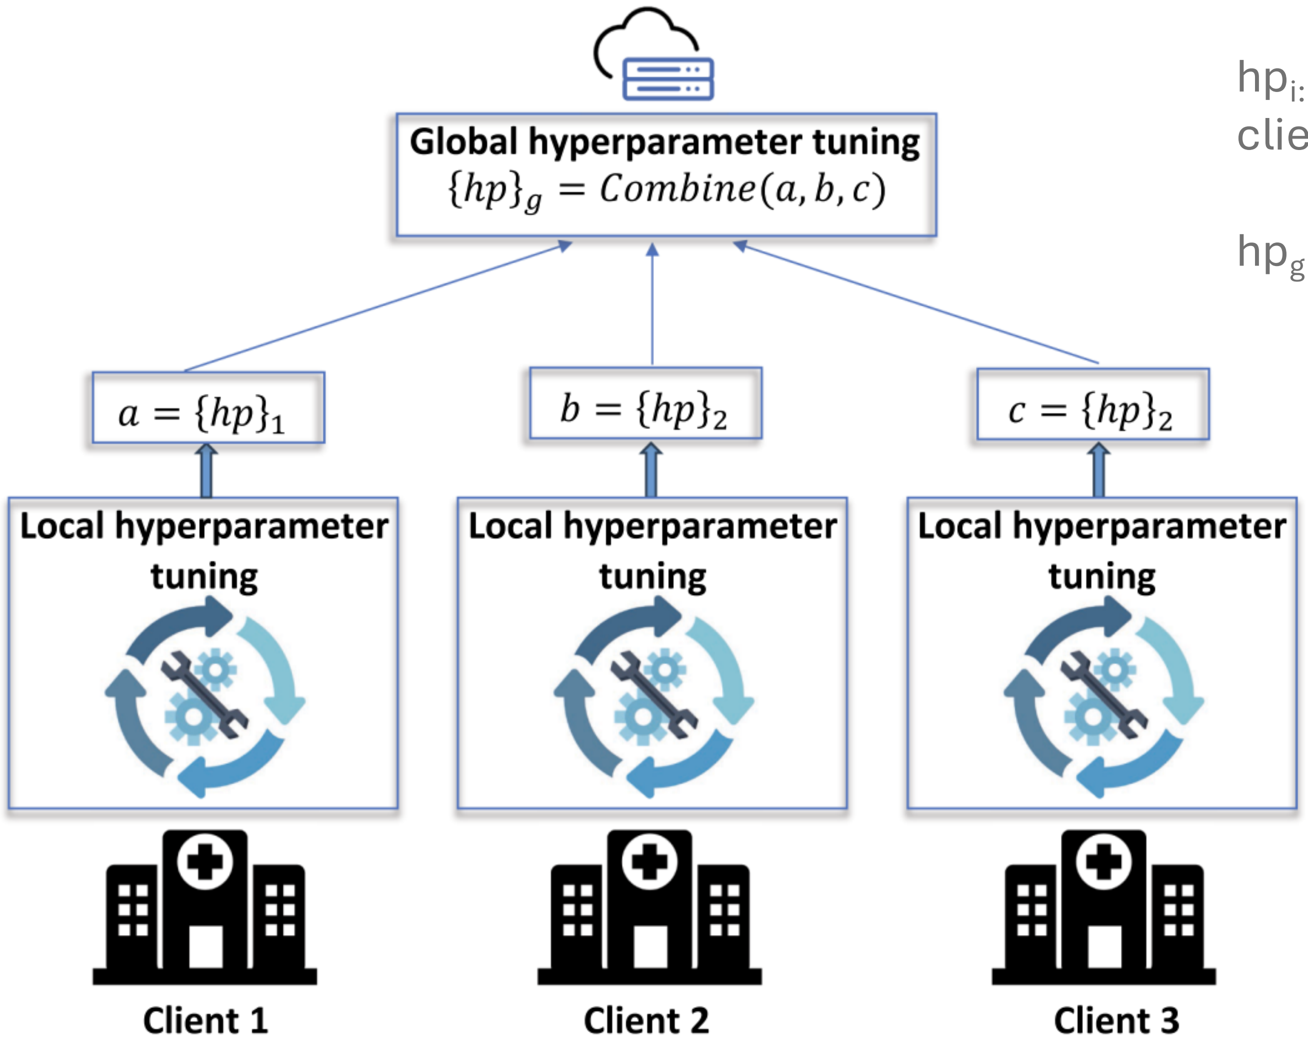

hp_i: best HP configuration at client i

hp_g: global best configuration

Research Questions

1) What is a suitable function to "Combine" the clients' hyperparameters and yield an effective global configuration?

Measurement Study

2) How to enable the "Combine" function while protecting the privacy of each client's hyperparameters?

PrivTuna Framework

Measurement Study

Methodology

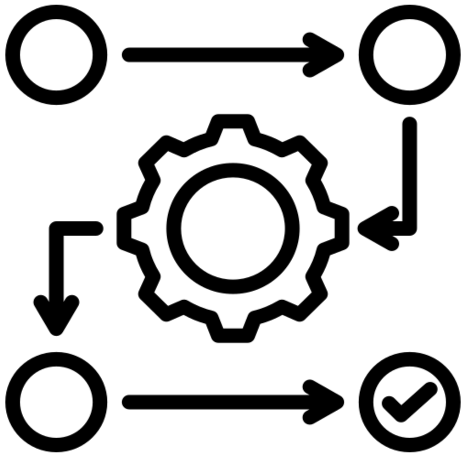

1) We perform local HP optimization (LHO) at each client to derive optimal HPs on local datasets
2) We perform global HP optimization (GHO) with federated grid search to construct a *ground truth* for the server HPs
3) We benchmark various *combination* strategies that yield the global HPs based on the local ones:
   - Mean
   - Median
   - Trimmed mean
   - Mean/Median of best HPs
   - Density based clustering
4) We compare the results of each HP combination strategy with GHO

Measurement Study

Experiment Setup

Datasets	Models	FL Settings	Hyperparameters
MNIST EMNIST Street View House Numbers CIFAR-10	1-layer CNN (30K params) 2-layer CNN (54K params) 6-layer CNN (700K params)	iid non-iid (label, feature, quantity skew)	Learning rate Momentum

We conducted over 4,067 experiments for the iid and 6,912 for the non-iid settings

Measurement Study

IID Setting Results

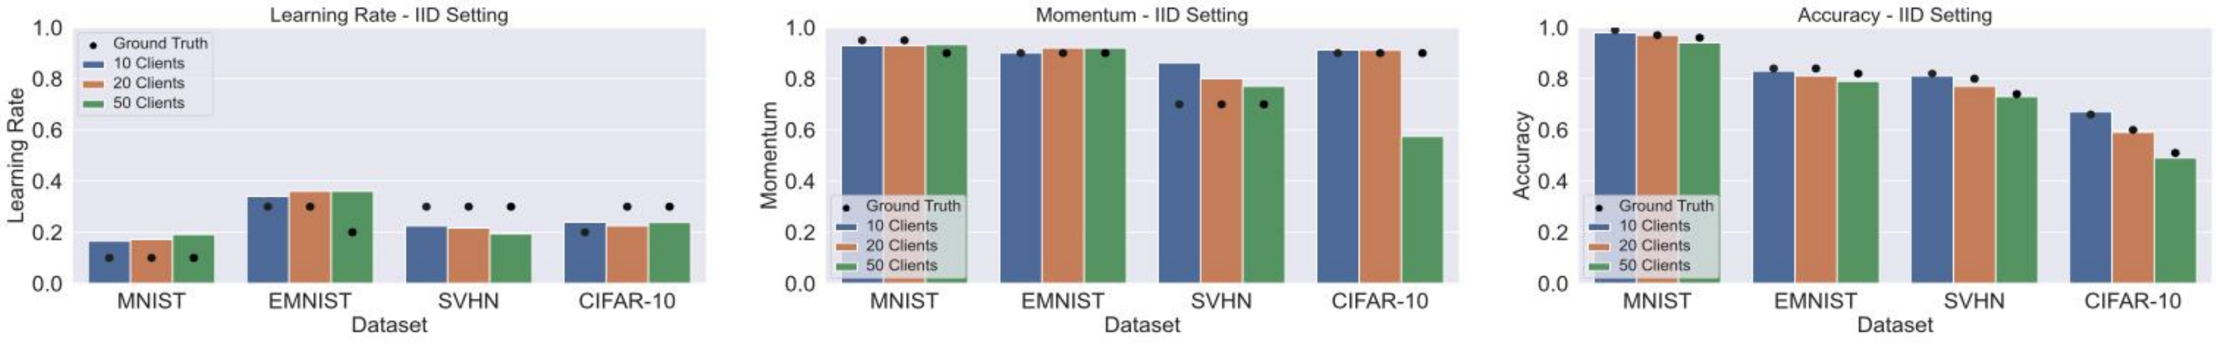

(a) Learning Rate (b) Momentum (c) Accuracy

HP averaging achieves good enough performance as the combination function

Measurement Study

Non-IID Setting Results

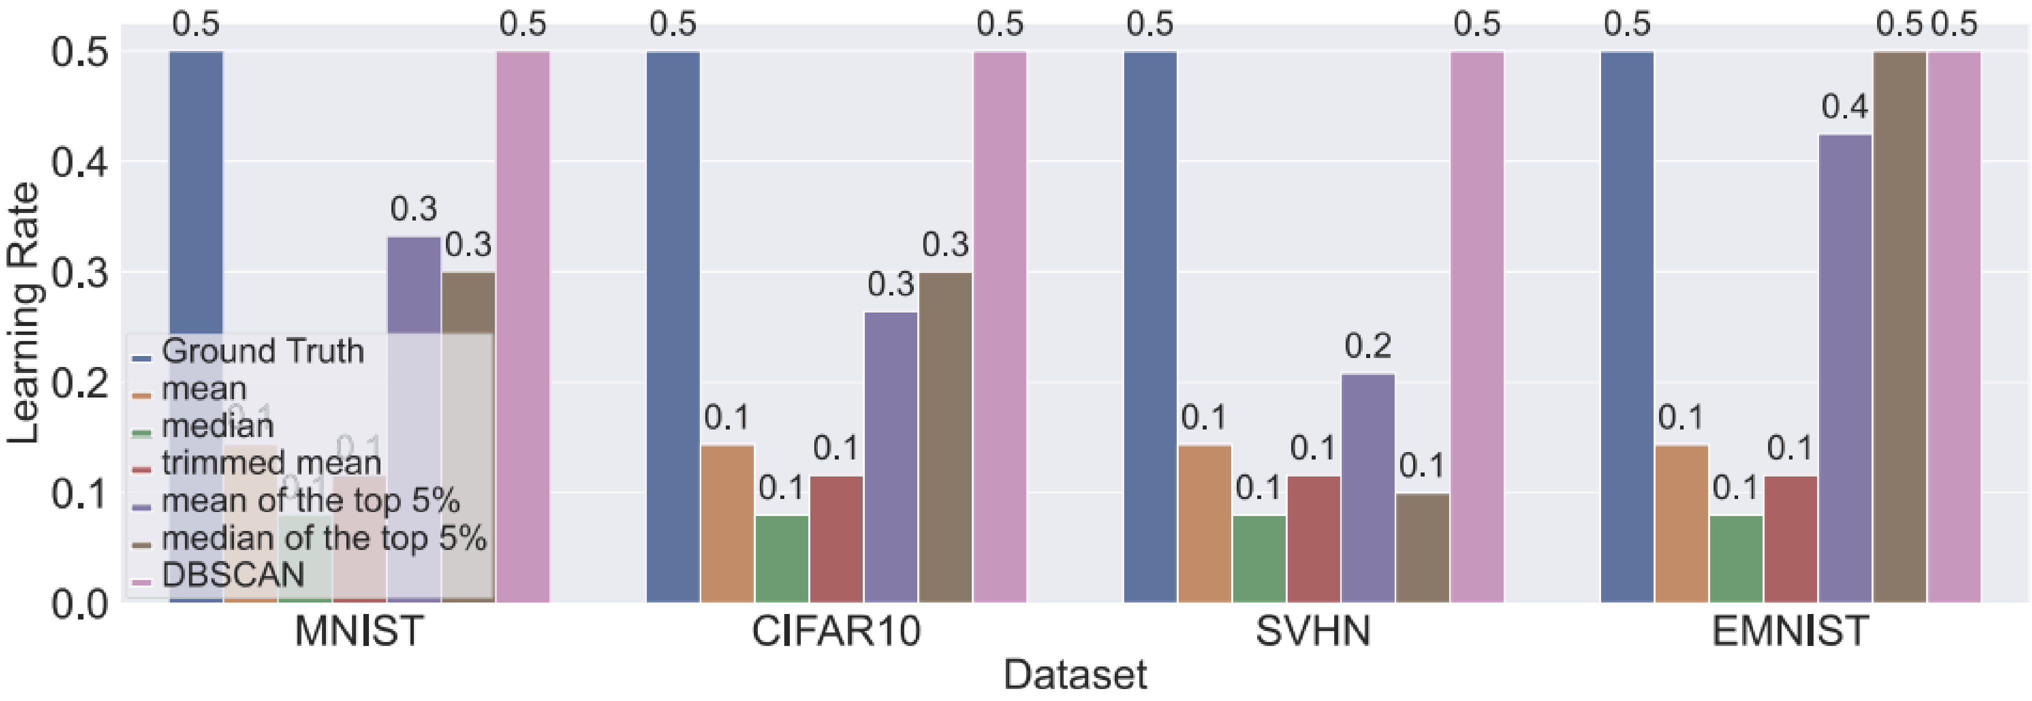

Learning rate with feature skew (β=0.02)

DBSCAN yields the best performance among various techniques (95/140 experiments)

Measurement Study

Comparison with Prior Work

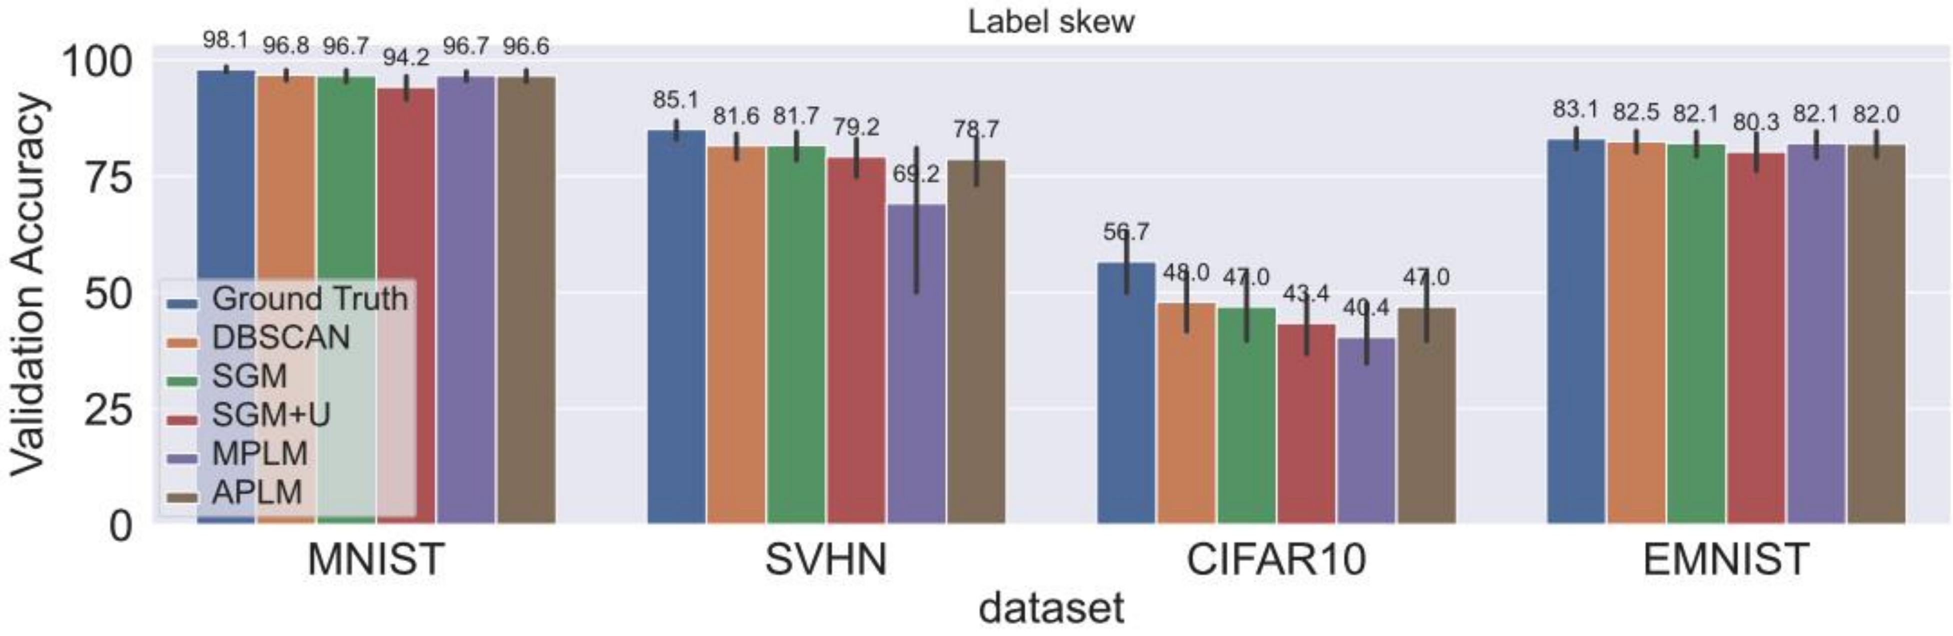

(c) Validation accuracies averaged across the label skew setting.

DBSCAN is on par with the SotA HP optimization method[1]

[1] Y. Zhou, et al., "Single-shot general hyper-parameter optimization for federated learning," in *ICLR*, 2023.

PrivTuna Framework

PrivTuna Framework

PrivTuna enables the "combine" function in an encrypted (MHE) and federated manner

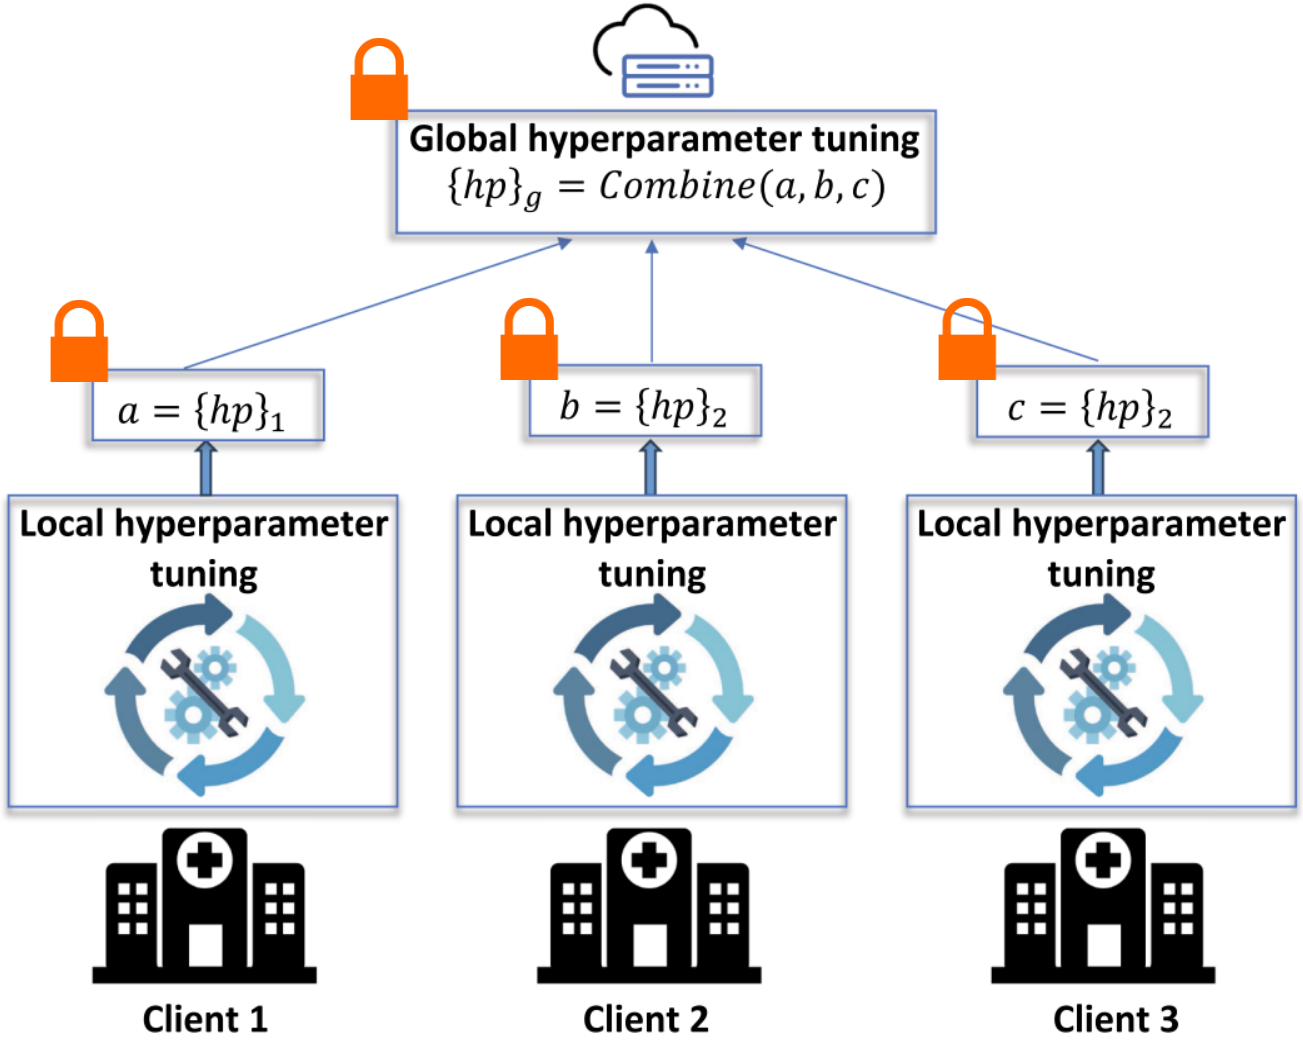

- MHE enables distributed trust (decryption requires collaboration between the clients)
- Balanced computation and communication overhead (e.g., plaintext/ciphertext operations, collective bootstrapping)

PrivTuna Overview

- **SecKeyGen(1^λ)**: Each client generates its private/public key pair (sk_i, pk_i)
- **DKeyGen({sk_i})**: The clients collectively generate a public key **pk** (and evaluation keys **evk**)
- Each client encrypts its LHO results with **pk**, i.e., **c_i = Enc(pk, LHO_i)**
- The server homomorphically executes the Combine functionality on the encrypted **c_i**
- **DDecrypt(c, {sk_i})**: The clients collectively decrypt the global hyperparameters

PrivTuna Framework

Experimental Results

Federated grid search requires ~2 hours of computation while POSEIDON[1] PPFL ~147 hours to tune HPs

	Runtime (s)	Comm. / Client (MB)	Precision (MSE)
PF-Mean	1.8	5.2	$1.8 * 10^{-4}$
PF-DBSCAN	17.3	28.2	$1.02 * 10^{-3}$

HW: Apple M2 Pro 3.49 GHz / 16 GB RAM

N: 20 clients

HE: CKKS

λ: 128-bit security

[1] Sav et al., POSEIDON: Privacy-Preserving Federated Neural Network Learning, *NDSS*, 2021.

Conclusion

- We investigated the problem of HP tuning in cross-silo FL
- We performed a comprehensive measurement study to understand the relationship between client and server HPs and identify strategies for single-shot HP tuning
- We introduced PrivTuna, an MHE framework for privately tuning HPs in cross-silo FL

The End

Thanks for your attention!

Full Paper: https://ieeexplore.ieee.org/document/10848179

Journal: IEEE Transactions on Privacy, Volume 2, 2025

Authors: Natalija Mitic[1], Apostolos Pyrgelis[2], Sinem Sav[3]

Affiliations: [1]Kera Health Platforms Inc, [2]RISE, [3]Bilkent University

Source Code: https://github.com/sinemsav/hyperparams

Contact: apostolos.pyrgelis@ri.se